Scalable Automated Detection of Spiral Galaxy Arm Segments

Darren R. Davis, Wayne B. Hayes
University of California, Irvine
Irvine, California 92697-3435 USA
drdavis@ics.uci.edu, wayne@ics.uci.edu

Given an approximately centered image of a spiral galaxy, we describe an entirely automated method that finds, centers, and sizes the galaxy and then automatically extracts structural information about the spiral arms. For each arm segment found, we list the pixels in that segment and perform a least-squares fit of a logarithmic spiral arc to the pixels in the segment. The algorithm takes about 1 minute per galaxy, and can easily be scaled using parallelism. We have run it on all ∼644,000 Sloan objects classified as "galaxy" and large enough to observe some structure. Our algorithm is stable in the sense that the statistics across a large sample of galaxies vary smoothly based on algorithmic parameters, although results for individual galaxies can sometimes vary in a non-smooth but easily understood manner. We find a very good correlation between our quantitative description of spiral structure and the qualitative description provided by humans via Galaxy Zoo. In addition, we find that pitch angle often varies significantly segment-to-segment in a single spiral galaxy, making it difficult to define "the" pitch angle for a single galaxy. Finally, we point out how complex arm structure (even of long arms) can lead to ambiguity in defining what an "arm" is, leading us to prefer the term "arm segments".

arXiv:1402.1910v1 [astro-ph.GA] 9 Feb 2014

I. INTRODUCTION

The Hubble Ultra Deep Field represents about 1/13,000,000 of the celestial sphere and contains about 10,000 galaxies, suggesting the entire sky contains upwards of 10^{11} galaxies. Gaining quantitative structural information for this number of galaxies will require automated methods. For spiral galaxies, existing methods for visual classification [1, 2] are either subjective or non-quantitative, while currently available semi-automated methods [3–11]. are either too simplistic or require significant human input.

II. METHODS

Our method is entirely automated, and is described in detail elsewhere [12]. Starting with an approximately centered image of a galaxy, it uses an iterative 2D Gaussian fit to find, precisely center, and size the image.

We make the simplifying assumption that the disk of the galaxy would be circular if viewed face-on, and then deproject it to produce a face-on view of the galaxy (Figure 1a). After applying a contrast-enhancement filter (Figure 1b), an orientation filter [9] is used to assign an orientation (strength and direction) to each pixel in the image based on the pixels around it (Figure 1c). Pixels are then hierarchically clustered into regions with locally similar orientation and consistent logarithmic spiral shape (Figure 1d). *Note that brightness plays no explicit role in this clustering, although it plays an implicit role through the creation of the orientation field.* Each cluster is associated with a logarithmic spiral arc determined by a least-squares fit, which can be brightness-weighted (also Figure 1d). Sometimes, the requirement for consistent logarithmic spiral structure blocks merges of clusters that, in retrospect, "should" have been merged. Later, however, as the clusters grow into their final shape, the arc fits may become more compatible than they first appeared. Thus, a second stage of merging is performed, based primarily upon compatible spiral arc parameters (Figure 1e). Figure 1f depicts the resulting arcs overlayed on the original de-projected image.

The resulting arcs and arm segments are independent of each other, do not need to conform to any symmetry criterion or be attached to a bar or the bulge, and do not need to all wind in the same direction. So far as we are aware, this list of arm segments comprises the most detailed description currently available for general spiral structure. Given the list of pixels for each segment, an astronomer could easily perform whatever measurement of that segment they may wish, such as color, luminosity, brightness profile, etc.

When run on a set of galaxies, our algorithm produces two tables: one table lists every arc in every galaxy, while the second offers a per-galaxy summary.

III. RESULTS

Figure 2 shows some typical "nice" examples. In general, we find that if the image is clear and has a sufficient signal-to-noise ratio (about as good as needed by a human, no more), then the algorithm does a very good job of marking out the arms and determinining pitch angle. However, we believe the strength of our algorithm is that it *also* does well when the galaxy is less clean, less symmetrical, and more fragmented. Even when the image has such low resolution that little structure is visible, we find good agreement with human determinations of structure.

Observing the arcs superimposed on the images in Figures 1, 2 and 3, we make two observations. First, that the logarithmic spiral arc is a very good mathematical description of the curve of spiral arms (with the possible exception of Figure 3d, which is a version of the code that is no longer in use and did not always properly handle multi-revolution arc fits). And second, that even though the fit is usually very good, the pitch angle between different arms in the same galaxy can be quite different. Even the two longest arms in the galaxy of Figure 1 do not agree with each other: the "left" long arm in Figure 1f has a pitch angle of 13.6 degrees while the "right" long one has a pitch angle of 24.4 degrees—a difference of almost 11 degrees. Furthermore, it is clear from Figure 1e that both

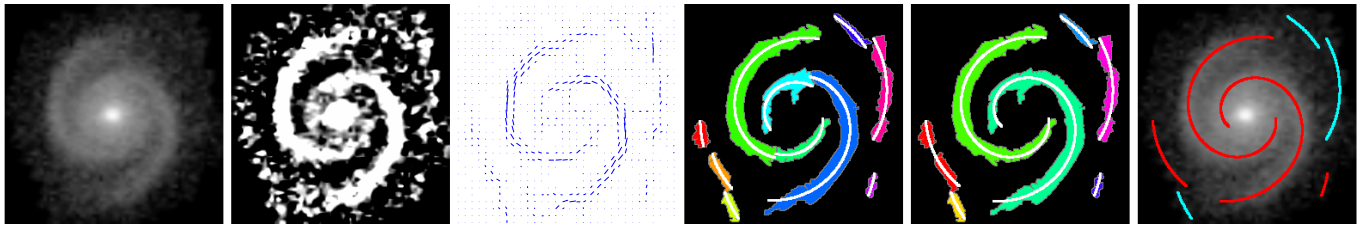

(a) (b) (c) (d) (e) (f)

FIG. 1: Steps in describing a spiral galaxy image. **a)** The centered and de-projected image. **b)** Contrast-enhanced image. **c)** Orientation field (at reduced resolution for pedagogical reasons). **d)** Initial arm segments found via Hierarchical Agglomerative Clustering (HAC) of nearby pixels with similar orientations and consistent logarithmic spiral shape, and the associated logarithmic spiral arcs fitted to these clusters. **e)** Final pixel clusters (and associated arcs) found by merging compatible arcs. **f)** Final arcs superimposed on image (a). Red arcs wind S-wise, cyan Z-wise.

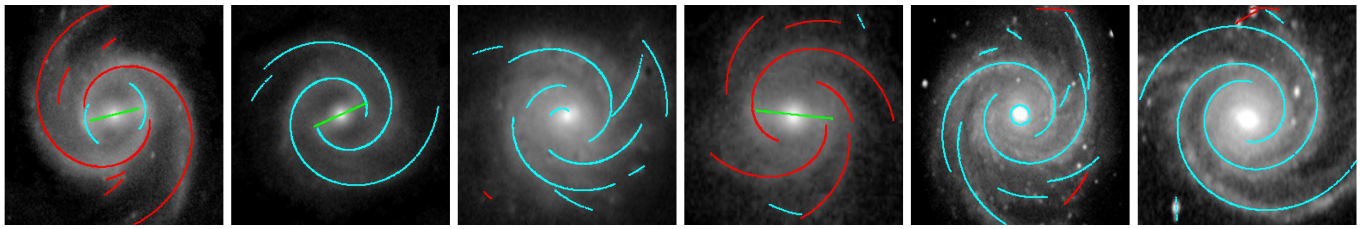

FIG. 2: Typical examples of how the algorithm performs on "nice" galaxies. Green represents a bar; red arms wind S-wise, cyan arms wind Z-wise.

arms are fit quite well, with a width much smaller than their length and the arcs remaining inside their cluster for their entire length; this indicates that the error in these pitch angles is much less than their difference. Such a disagreement is not atypical, as Table I shows. Recent work has suggested that pitch angle can also vary with radius, and this variance correlates with other structural parameters of the galaxy [13].

minimum length	0	50	100	150	200	250	300
median difference	14.5	14.3	10.7	7.5	5.6	3.5	2.6

TABLE I: Median pitch angle difference (bottom row, in degrees) between the two longest arms in a galaxy when both are at least as long (in pixels) as the top row. Images are scaled to 256x256 pixels.

Astronomers frequently refer to "spiral arms" in a galaxy, but to our knowledge there does not exist a formal definition of what, exactly, a spiral arm is. Figure 3 illustrates why this may be the case: arm segments can fork, split, have dust lanes, and have "kinks" in them where the pitch angle changes even if they are visually contiguous. As a result, in this paper we prefer to use the term "arm segments", which we roughly define as an (almost) contiguous arc-shaped region of light with locally consistent orientation. Figure 3 also gives an informal indication of how the parsing of an individual image can change depending upon minor changes to the algorithm.

Given a fixed version of the algorithm in which only its major parameters are changed, it is important to determine that our results change in a reasonable way as a function of the parameters. Although the results for an individual galaxy may occasionally change abruptly as algorithmic parameters change—especially if the galaxy is poorly resolved—the average results across a large sample of galaxies should change smoothly. We have performed a sensitivity analysis and verified that, statistically, all output variables vary smoothly. Furthermore, both the mean and median pitch angle across galaxies changes by less than a degree as algorithmic parameters are varied by a factor of 2 in either direction. This gives us confidence that the results of our algorithm are stable in the sense that statistics across a large sample of galaxies are unlikely to suffer a large or unexpected shift based on small changes to the parameters.

So far as we are aware, the current work is the first large-scale, automated, quantitative survey of detailed general spiral structure in galaxies. There exist several small-scale quantitative surveys (*e.g.*, [7, 8]), and two large-scale, qualitative, human-based surveys: one citizen science project (Galaxy Zoo [14, 15]), and one professional survey [16].

We compare our algorithm to two small-scale pitch angle surveys. In the first Ma [8] manually measured the pitch angles of either one or two arms in each of a small sample of galaxy images. We have downloaded and run our algorithm on the images used in that paper. The second (a group from the University of Arkansas) uses a Fourier Analysis to extract the dominant pitch angle [7, 17]. In both cases, we always agree in chirality, although there is significant scatter, sometimes up to 10–20 degrees, when it comes to pitch angle agreement. Although this scatter may seem to be a concern, recall from Table I that the pitch angle can vary quite significantly between arms in a single spiral galaxy. Since Ma [8] measured only one or two arms in each galaxy, and since the Arkansas group measure only one dominant pitch angle, perhaps the scatter could be explained by inter-arm differences in

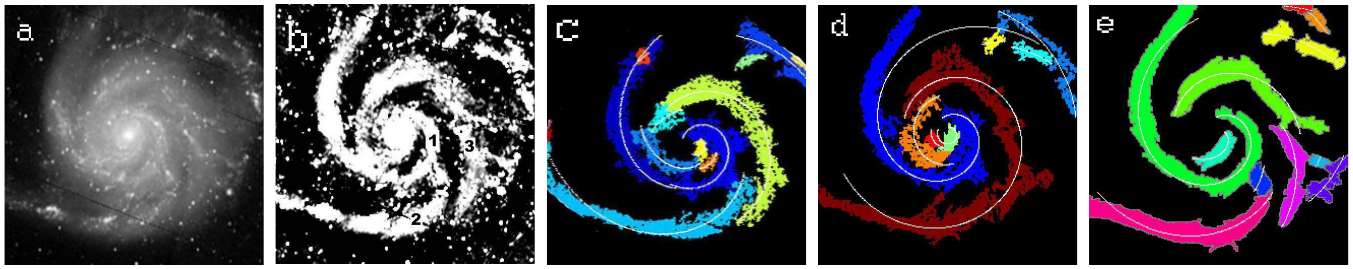

FIG. 3: Three different interpretations of the spiral structure in M101. The colored images are from different versions of the code. All comprise reasonable interpretations of the structure. **(a)** original image **(b)** Contrast enhanced. We have labelled three arm segments; the joint between segments 1 and 2 may be what a human would call a "fork", although our code never refers to forks. **(c)** An old version of the code, where the three segments happen to be separated (blue, cyan, and olive pixel clusters). **(d)** An intermediate version of the code, where segment 2 has been interpreted as a continuation of segment 3, jumping over the gap between them. The single logarithmic spiral arc spanning the two arguably fits reasonably well, suggesting perhaps that segments 2 and 3 are physically one arm with an obscuring dust lane, while the apparent "fork" of segment 2 from segment 1 is an optical illusion. **(e)** The most recent version of the code, in which the logarithmic spiral arc more stingently fits (in the least-squares sense) each cluster of pixels.

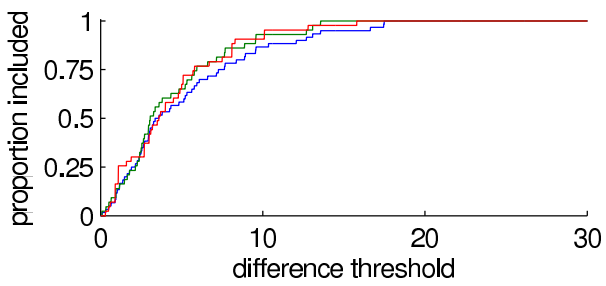

FIG. 4: Cumulative distribution of pitch angle discrepancies between the two arcs measured for each galaxy from [8], in red, and between our measurements and the measurements in [8] (blue for the full comparison set, and green for the subset where two arc measurements are available from [8]). Since the two curves are very similar, it is likely that much of the scatter between our results and the results of other authors arises from genuine within-galaxy arm variation rather than between-method measurement variation.

each galaxy, rather than differences as a function of method. To test this hypothesis, Figure 4 plots the cumulative distribution of pitch angle discrepancies between the two arcs in one galaxy (when available) from Ma [8], *vs.* discrepancies between Ma's method and our method.

Galaxy Zoo [14, 15] is a citizen science project in which approximately 250,000 human volunteers classify images of galaxies over the web after some rudimentary training. The median number of people who viewed each image was about 40, so that some measure of certainty can be obtained from multiple viewings of each image. Galaxy Zoo 1 (GZ1, [15]) presented people with an image of a real galaxy along with 6 cartoon galaxies, and asked them to choose which of them most resembled the real galaxy. Across spiral galaxies with observable structure, the only comparison we can make with GZ1 is chirality (S-wise vs. Z-wise). In difficult cases some humans may choose "spiral" while others do not, so we compare our results on chirality against what we call the "discernability" of a galaxy: the maximum fraction of agreeing humans that saw spiral structure, which we define for a particular galaxy as $\frac{\max(\text{S-wise votes, Z-wise votes})}{\text{total number of votes}}$.

min discernibility rate	0	60	90	100	0	60	90	100
require 2 longest agree	N	N	N	N	Y	Y	Y	Y
inclusion rate	99.3	95.1	52.2	12.4	67.0	64.7	39.0	10.0
per-arc majority vote	75.4	75.9	79.9	82.5	79.5	79.8	83.1	84.8
longest arc alone	84.9	85.3	89.4	92.4	95.3	95.7	97.8	98.4
length-weighted vote	89.4	89.9	93.5	95.7	94.9	95.3	97.5	98.4

TABLE II: Winding-direction agreement with human classifications from Galaxy Zoo 1. Row 1: the minimum proportion of human votes that the dominant winding direction must receive (out of the two known-direction and four unknown-direction categories). Row 2: do we demand our two longest arcs agree in chirality? Row 3 gives the proportion of the 29,250 galaxies included under these criteria. Rows 4, 5, and 6 give agreement rates between Galaxy Zoo and three methods of determining winding direction from our output.

Table II compares our chirality measurement against those of Galaxy Zoo 1 humans as a function of both human discernability, and several measures of our output. We include 29,250 galaxies with clearly observable structure, chosen by the director of the GZ project (Steven Bamford, personal communication). We see that both the "Longest arc alone" and "Length-weighted vote" agree with the humans at least 85% of the time, and as much as 98.4% of the time, across all values of human discernability.

We also compare pitch angle (arm tightness) measurements, starting with Galaxy Zoo 2. Here, where human classifiers indicated that there was "any sign of a spiral arm pattern" [18], they were asked whether the spiral arms were tight, medium, or loose. Figure 5 shows the relationship between our measured pitch angle (arclength weighted vote) and the proportion of galaxies receiving a majority human vote for Tight, Medium, or Loose. As can be seen from the dashed lines in Figure 5, galaxies where we measure a low pitch angle usually have majority human votes for Tight, while most of the remaining galaxies in this range had majority votes for Medium. As our measured pitch angle increases, we see pro-

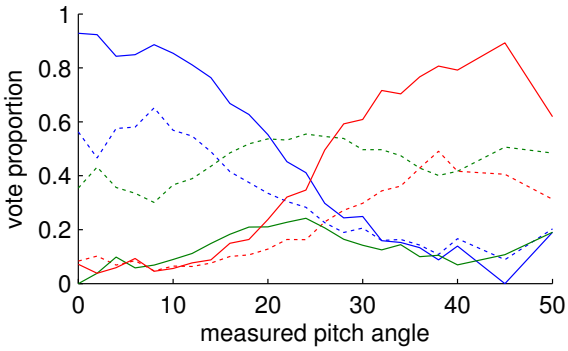

FIG. 5: Proportion of galaxies receiving a majority vote for Tight (blue lines), Medium (green lines) or Loose (red lines) as a function of our measured pitch angle. Pitch angles are binned with width 2 degrees between 0 and 40, 5 degrees between 40 and 50, and one bin beyond 50 (due to low sample size). The dashed lines represent all images tested from Galaxy Zoo 2; the solid lines represent the images within the top quartile of human agreement.

gressively fewer galaxies classified as Tight, and more galaxies classified as Loose. Designations as "Medium" are pervasive throughout, while "Loose" classifications are less frequent than "Tight" classifications. This reflects the classification distribution of the image set as a whole. The solid lines in Figure 5 represent the top agreement quartile (lowest Shannon entropy quartile), and the association between our tightness measure and human classifications is even more pronounced.

Longo [19] also provided a chirality measurement for a survey of galaxies, in which each galaxy was viewed only once by one of five student volunteers. The students were told to choose a chirality only if it was clear. We find similar agreement as was seen against Galaxy Zoo.

IV. DISCUSSION AND FUTURE WORK

We have run our code on every item in the Sloan Digital Sky Survey (SDSS) that is classified as "galaxy". Unfortunately, SDSS does not distinguish between spiral and non-spiral galaxies. We are currently working on a machine learning algorithm that uses the output of our code to distinguish between spiral and non-spiral galaxies. Preliminary results are encouraging (agreeing approximately 90–95% with Galaxy Zoo humans), and will be presented in a future paper. Once we can separate out spiral galaxies, further studies will be performed concerning how spiral structure correlates with other variables such as color, redshift, local environment, etc.

Code and data are available from WBH upon request.

Acknowledgments

We thank Steven Bamford for helpful insights, image sample selection, and pre-publication access to Galaxy Zoo 2 classifications; Charless Fowlkes, Deva Ramanan, Eric Mjolsness, and Aaron Barth for helpful discussions; and the Arkansas Galaxy Evolution Survey (AGES) Collaboration for pitch angle measurements and discussions. Comparisons were also made possible due to image data from CGS as well as from SDSS and POSS II. Fellowship and travel support was provided by an ICS Dean's Fellowship at UC Irvine (for DD); the Oxford Centre for Collaborative Applied Mathematics; Steven Bamford and the MegaMorph project; and the AGES Collaboration (through NASA Grant NNX08AW03A).

[1] E. P. Hubble, *The Realm of the Nebulae* (Yale University Press, 1936).
[2] G. de Vaucouleurs, Handbuch der Physik **53**, 275 (1959).
[3] R. E. de Souza, D. A. Gadotti, and S. dos Anjos, Ap.J. Suppl. S. **153**, 411 (2004), ISSN 0067-0049, .
[4] L. Simard, in *Astronomical Data Analysis Software and Systems VII* (1998), vol. 145, p. 108, ISSN 1050-3390, .
[5] C. Y. Peng et al., Astron. J. **124**, 266 (2002), ISSN 00046256, .
[6] C. Y. Peng et al., Astron. J. **139**, 2097 (2010), .
[7] B. L. Davis et al., Ap.J. Suppl. S. (2012), accepted, 1202.4780.
[8] J. Ma, Chin. J. Astron. Astrophys. **1**, 395 (2001), .
[9] K. Au, Phd thesis, Carnegie Mellon Univ. (2006).
[10] B. D. Ripley and A. I. Sutherland, Philos. Trans. R. Soc. London, Ser. A **332**, 477 (1990), .
[11] L. Shamir, Ap.J. **736** (2011), ISSN 0004-637X, .
[12] D. Davis and W. Hayes, in *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on* (2012), pp. 1138–1145, ISSN 1063-6919.
[13] S. S. Savchenko and V. P. Reshetnikov, Monthly Notices of the Royal Astronomical Society **436**, 1074 (2013).
[14] C. J. Lintott et al., MNRAS **389**, 1179 (2008), .
[15] C. Lintott et al., MNRAS **14**, 1 (2010), arXiv:1007.3265v1, .
[16] P. B. Nair and R. G. Abraham, Ap.J. Suppl. S. **186**, 427 (2010), ISSN 0067-0049, .
[17] M. S. Seigar et al., Ap.J. **645**, 1012 (2006), ISSN 0004-637X, .
[18] K. L. Masters et al., MNRAS **411**, 2026 (2011), ISSN 00358711, .
[19] M. J. Longo, Phys. Lett. B **699**, 224 (2011), ISSN 03702693, .